



Bill White
Rav Meerwald
Jeff Nesbitt

zSeries Qualified WDM Vendor: Lucent Technologies

Introduction

This IBM® Redpaper is one in a series describing zSeries® qualified optical Wavelength Division Multiplexing (WDM) vendor products for Geographically Dispersed Parallel Sysplex™ (GDPS®) solutions.

International Business Machines (IBM) Corporation and Lucent Technologies have successfully completed testing of the Lucent Technologies Metropolis® Enhanced Optical Network (EON) System. The IBM Vendor Solutions Connectivity Lab has accepted the test results for inclusion in the zSeries GDPS Qualification program. This Redpaper describes the applicable environments, protocols, and topologies qualified and supported by IBM @server zSeries for connecting to the Lucent Metropolis® EON product hardware and software at release level 8.2.1. This Redpaper version supersedes and replaces all previous versions. The latest version of this Redpaper is available at:

<http://www.redbooks.ibm.com/abstracts/redp3906.html?open>

IBM @server zSeries GDPS Qualification and IBM TotalStorage® Proven™ are complementary programs leveraging the strengths of IBM licensed technologies, intellectual property, and know-how, enhancing interoperability in a multivendor environment. The two programs continue to work together as new technologies and products become available. The Lucent Technologies Metropolis® EON System is also IBM TotalStorage Proven.

The zSeries qualified WDM vendor products that support GDPS are:

- ▶ Adva Fiber Service Platform 2000 (FSP 2000) DWDM system
- ▶ Cisco ONS 15530 DWDM, ONS 15540 ESPx, and ONS 15454 MSTP
- ▶ IBM 2029 Fiber Saver DWDM (withdrawn from marketing)
- ▶ Lucent Technologies Metropolis® Enhanced Optical Network (EON) System
- ▶ Nortel Networks OPTera Metro 5000 Multiservice Platform series

In addition to the above WDM products, the IBM 9036 Model 003 extender is also available to extend Sysplex Timer® (ETR and CLO) links from 3 km to 26 km. An IBM 9036-003 is available as a Request for Price Quotation (RPQ 8K1919).

zSeries GDPS qualification overview

Geographically Dispersed Parallel Sysplex™ (GDPS)®, an industry-leading ebusiness continuity solution, is a multisite solution designed to provide the capability to manage the remote copy configuration and storage subsystems, automate Parallel Sysplex operational tasks, and perform failure recovery from a single point of control, thereby helping to improve application availability. GDPS supports both the IBM TotalStorage Metro Mirror (synchronous Peer-to-Peer Remote Copy (PPRC)), as well as the IBM TotalStorage z/OS® Global Mirror (asynchronous Extended Remote Copy (XRC)) forms of remote copy. Depending on the form of remote copy implemented, the solution is referred to as GDPS/PPRC or GDPS/XRC.

IBM only supports WDM products qualified by IBM @server zSeries for usage in GDPS/PPRC and GDPS/XRC solutions. To obtain this qualification, WDM companies obtain licensed IBM patents, intellectual property, and know-how related to the GDPS architecture. This gives them access to the unique IBM protocols and applications used in a GDPS environment (including Sysplex Timer, InterSystem Channel (ISC), PPRC, and XRC). Without this information, it is not possible to reverse engineer the requirements for a GDPS environment. Qualified companies will typically license this information for an extended period of time, allowing them to subscribe to the latest GDPS architecture changes, and be among the first to market with offerings that support these features.

In addition, these vendor products have been tested and qualified by IBM with the same test labs and procedures used to test all aspects of a GDPS environment. This testing includes functionality, recovery, and in some cases performance measurements. Having access to these test facilities allows IBM to configure a fully functional sysplex, and simulate failure and recovery actions, which could not be tested as part of a working client environment. IBM has the facilities to test and qualify these products with both current and legacy equipment, leveraging the size of the Vendor Solutions Connectivity Lab in Poughkeepsie, USA. This qualification testing allows IBM to reproduce any concerns that may arise while using this equipment in a client's application. Figure 1 on page 2 shows a GDPS/PPRC environment used for WDM vendor qualification.

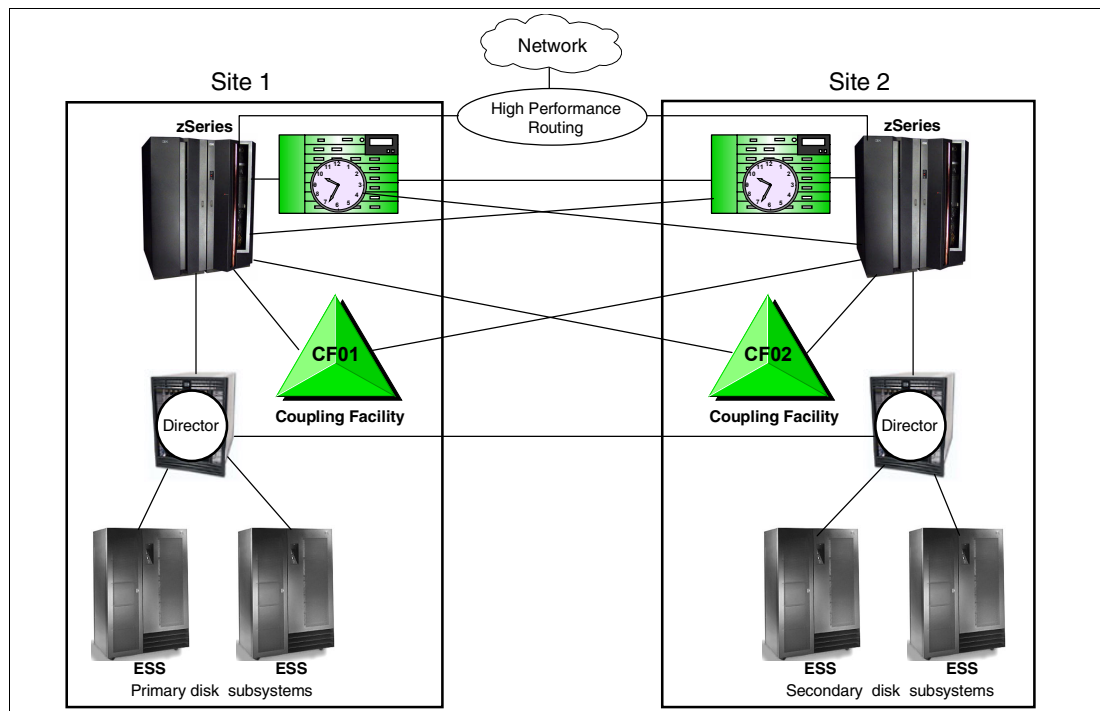


Figure 1 GDPS environment

Qualification testing

GDPS qualification testing is carried out at the IBM Vendor Solutions Connectivity Lab in Poughkeepsie, New York. The Vendor Solutions Connectivity Lab contains zSeries hardware systems with software applications allowing complete functional testing of a GDPS environment, as well as testing the interoperability of WDM products within a GDPS.

Hardware used for testing purposes include (but is not limited to) the following:

- ▶ IBM @server zSeries servers
- ▶ IBM 9037 Model 2 Sysplex Timers
- ▶ IBM 2105 TotalStorage Enterprise Storage Servers
- ▶ IBM 9032 ESCON® Directors
- ▶ IBM @server zSeries qualified FICON® Directors and switches

WDM links of varying distances are deployed, using spools of single-mode fiber in lengths from 5 km to 50 km. Multiple spools are interconnected to test WDM link protocols up to the maximum qualified distance of 100 km (62 miles). To achieve the maximum distances qualified for GDPS protocols, vendors may utilize erbium doped fiber amplifiers (EDFAs), inserted at various link points, to condition the signals on the fiber link connecting the WDM equipment.

Operating system and application software is installed to create, and stress test, the GDPS environment. Software used in the test environment include (but is not limited to) the following:

- ▶ z/OS, Linux® for zSeries, and Parallel Sysplex software exploiters
- ▶ Coupling Facility Control Code (CFCC)
- ▶ IBM proprietary software and microcode utility test suites

As part of the GDPS qualification test, IBM proprietary software and microcode utility test suites drive the various GDPS components and protocols to the full data rate of each link type being transported by the WDM equipment. This ensures maximum channel throughput is achieved and tested to levels well beyond typical client environments.

The test suites are used for verification of z/Architecture functionality, with individual tests typically being run for up to 24 hours at a time. During these functionality tests, it is expected that zero errors will be detected by the attached subsystems for a test to be classified as successful. Any errors detected during this testing are captured and analyzed by the test suites.

The test suites are also used for verification of z/Architecture recovery by creating various fault and error conditions. The recovery tests check for the correct detection of a fault or error condition by the attached subsystems, and ensure that the recovery adheres to z/Architecture rules.

Some of the recovery tests performed for each link type include:

- ▶ Link state change interrupt detection and recovery. Links are deliberately broken and re-established to ensure that detection and error recovery takes place correctly.
- ▶ Link error threshold and link synchronization error, detection, and recovery. Errors are deliberately injected, at the application and channel subsystem levels, into the link protocol data streams to ensure that detection and error recovery takes place correctly.
- ▶ Link service and maintenance package recovery. Link hardware maintenance actions are performed to ensure that link state change detection and recovery takes place correctly.
- ▶ Testing of the various WDM protection schemes are performed to ensure that the expected link errors are detected, and recovery takes place correctly.

GDPS components and protocols

The following IBM technologies are functional components of GDPS and are tested during the qualification process. These components may also be used by clients in environments that do not require a full GDPS solution. The testing provides a level of assurance that these components will function in a WDM environment.

Components

The following GDPS components are tested during the qualification process:

- ▶ IBM Parallel Sysplex
- ▶ IBM TotalStorage® Enterprise Storage Server® (ESS)
- ▶ Optical Dense or Coarse Wavelength Division Multiplexer (DWDM or CWDM)
- ▶ IBM TotalStorage® Metro Mirror (Peer-to-Peer Remote Copy (PPRC)), a synchronous form of remote copy
- ▶ IBM TotalStorage® z/OS Global Mirror (Extended Remote Copy (XRC)), an asynchronous form of remote copy

The GDPS solution is also independent of disk vendor, as long as the vendor meets the specific levels of Metro Mirror (PPRC) and z/OS Global Mirror (XRC) architectures.

Protocols

The following GDPS connectivity protocols are tested during the qualification process:

- ▶ Enterprise Systems Connection (ESCON)
- ▶ IBM 9037 Sysplex Timer (ETR/CLO)
- ▶ Fiber Connection (FICON) (1Gbps)
- ▶ Fibre Channel (FC 100) (1 Gbps)
- ▶ Gigabit Ethernet
- ▶ InterSystem Channel (ISC) and ISC-2 (1 Gbps)
- ▶ InterSystem Channel (ISC-3) Compatibility Mode (1Gbps)
- ▶ InterSystem Channel (ISC-3) Peer Mode (2 Gbps)
- ▶ InterSystem Channel (ISC-3) Peer Mode (1 Gbps), via RPQ 8P2197

Table 1 lists the data transfer rates and maximum qualified distances for GDPS qualified protocols. For some extended distances, the use of optical amplifiers may be required.

Table 1 Qualified protocols and distances

Protocol	Data transfer rate	Max Repeated Distance
ESCON	200 Mbps	100 km ^{a b}
CLO	8 Mbps	40 km ^a
ETR	8 Mbps	40 km ^a
FICON (1 Gbps) ^c	1.062 Gbps	100 km
Fibre Channel FC100 (1 Gbps)	1.062 Gbps	100 km
Gigabit Ethernet	1.25 Gbps	100 km
ISC / ISC-2	1.062 Gbps	40 km
ISC-3 Compatibility Mode	1.062 Gbps	40 km

Protocol	Data transfer rate	Max Repeated Distance
ISC-3 Peer Mode	2.125 Gbps	100 km ^a
ISC-3 Peer Mode (1Gbps) ^d	1.062 Gbps	40 km

a. Requires RPQ 8P2263 zSeries Extended Distance (8P2262 for S/390®).

b. Effective channel data rate of an ESCON channel is affected by distance.

c. Including FICON Bridge card.

d. Requires RPQ 8P2197. This RPQ provides an ISC-3 Daughter Card which clocks at 1.062 Gbps in Peer and Compatibility modes.

IBM 9037 Sysplex Timer

The IBM 9037 Sysplex Timer is a mandatory component of GDPS/PPRC. The Sysplex Timer provides an External Time Reference (ETR) to synchronize the time of day (TOD) clocks on attached zSeries servers in a GDPS/PPRC environment.

The 9037 Sysplex Timer uses two link types:

- ▶ External Time Reference (ETR)

ETR links are connections between the Sysplex Timer and the zSeries server ETR ports providing TOD clock synchronization between multiple servers.

- ▶ Control Link Oscillator (CLO)

CLO links are connections between two Sysplex Timer units in High Availability mode allowing synchronization of the Sysplex Timer timing signals.

To ensure correct Sysplex Timer and server time synchronization, the end-to-end lengths of the transmit and receive fibers within an individual ETR or CLO link must be equal (within 10 meters).

Special care should be taken when using EDFA optical amplifiers or dispersion compensation modules (DCMs) to ensure the end-to-end lengths of the transmit and receive fibers of the link are equal (within 10 meters). EDFAs and DCMs contain significant lengths of fiber, which must be included in the total fiber distance calculation.

Technical description

The Lucent Technologies Metropolis® EON System is a scalable, high-speed fiber-optic data transport system. It consists of multiple bays that can be interconnected to form an optical network supporting up to 32 International Telecommunications Union (ITU) wavelengths. These ITU wavelengths are multiplexed onto a single pair of fibers using optical wavelength division multiplexing (WDM).

The optical network connecting the shelves can be configured in a two site point-to-point or a multi-site ring network. A fully configured point-to-point network can support up to 32 client interfaces transported by up to 32 protected ITU wavelengths over a pair of fibers (trunks) connected between data sites.

The zSeries qualified product in the Lucent Technologies Metropolis® Enhanced Optical Network (EON) System is:

- ▶ Lucent Metro Enhanced Optical Network (EON), Release 8.2

The Lucent Metro EON consists of equipment bays or racks that house several card shelves. Two equipment bays can be interconnected to support up to 32 protected client interfaces transported over 32 ITU wavelengths.

Interface cards and modules

The following is a detailed list of the Lucent Metro EON optical interface cards and circuit packs qualified by IBM @server zSeries:

► Optical Translator Unit (OTU)

The Optical Translator Unit (OTU) card connects client equipment to the WDM platform. The OTU card converts client optical signals to an ITU compliant wavelength for transmission to the remote site.

The qualified OTU card types are:

- Low-Speed Broadband (LSBB), 8 Mbps to 750 Mbps, 1310 nm (1 port, fixed transceiver)
- Enhanced Low-Speed Broadband (ELSBB) 8 Mbps to 1.062 Gbps Mbps, 1310 nm (1 port, fixed transceiver)
- Universal Broadband (UBB), 8 Mbps to 2.5 Gbps, 1310 nm (1 port, fixed transceiver)

► Optical Unidirectional Path Switching Ring (OUPSR)

The OUPSR circuit pack provides per wavelength protection for client interfaces. The OUPSR circuit pack protects against failure of an OTU card, or from physical damage to a site-to-site fiber pair by switching to a redundant fiber pair

► Optical Amplifier (OA)

OA circuit packs provide in-line optical amplification of up to 32 ITU specific wavelengths within a single fiber. For some extended distances, the use of these erbium doped fiber amplifiers (EDFAs) may be required.

The qualified OA circuit packs are:

- LEA 207, Single stage, C-band
- LEA 307, Dual stage, C-Band

Topologies and protection schemes

IBM @server zSeries qualifies the two site point-to-point WDM network topology and protection against failures in site-to-site fiber links or failures in individual components within the WDM network for GDPS.

GDPS is a high availability solution that can utilize several protection schemes; some restrictions apply for particular protocols, for example, Sysplex Timer (ETR/ECLO) links.

Important: Protection schemes that cannot guarantee ETR and CLO link (transmit and receive) lengths within the 10 meter limit are not supported by the IBM 9037 Sysplex Timer.

Protection schemes

The Lucent Metro EON system provides the following protection schemes:

► Unprotected

An unprotected Optical Translator Unit (OTU) is connected to one client interface and to one site-to-site fiber link only. A failure of the OTU or the site-to-site fiber link will result in a loss of client communications.

► Client based protection

Client based protection uses at least two client interfaces connected to the WDM. These interfaces are arranged so that the OTU circuit packs connecting the two sites are distributed over two diverse site-to-site fiber links. The client device is responsible for ensuring that a failure of a WDM module or of a single site-to-site fiber link will not result in total loss of client communications.

IBM 9037 Sysplex Timer CLO links are qualified for use in a Client based protection scheme, as long as they are using separate paths (routes).

Note: For simplicity, not all components in the optical path are shown in the following diagrams.

Figure 2 on page 7 shows a high level view of the Client based protection scheme. In this case, a client device would have two separate site-to-site connections: one via OTU A and the other via OTU B.

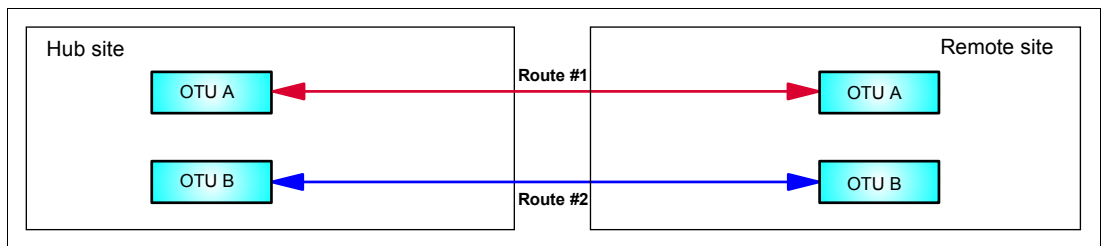


Figure 2 Client based protected scheme

► OUPSR protection

Optical Unidirectional Path Switched Ring (OUPSR) protection provides per wavelength protection against a site-to-site fiber failure, and failure of OTU circuit packs. In the transmit direction, the signal from a client is optically split and carried over redundant OTU circuit packs and diverse site-to-site fiber links. At the remote site, WDM control logic selects the signal from one of the OTU circuit packs as the active and the other as backup, for connection to the remote client interface. OUPSR protection is unidirectional, which means that if one fiber of a pair is damaged, only that fiber will be switched to the backup.

IBM 9037 Sysplex Timer ETR and CLO links cannot be used with OUPSR protected OTU circuit packs because OUPSR protection is unidirectional.

ISC, ISC-2 and ISC-3 Compatibility Mode links should not be used with OUPSR protection. Use Unprotected or Client Protection schemes for these link types

The Lucent Metro EON system implements a software lockout function to inhibit ETR, CLO, ISC, ISC-2 and ISC-3 Compatibility Mode links from being provisioned as protected.

Figure 3 shows a high level view of the OUPSR protection scheme.

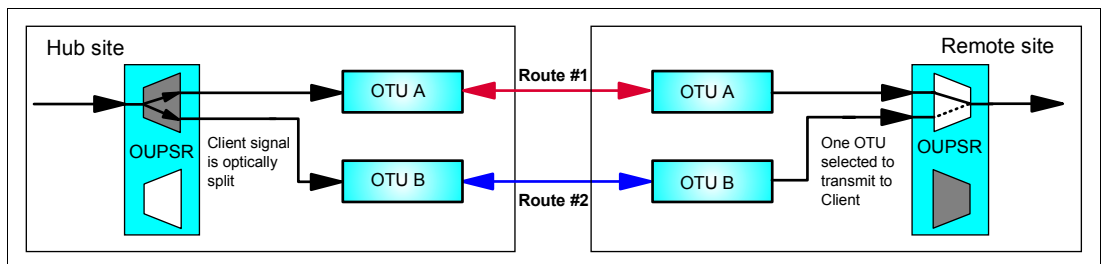


Figure 3 OUPSR protected scheme

Protection scheme intermix

Unprotected, Client protected, and OUPSR protected schemes can be intermixed within the same Lucent Metro EON system or network on a per client interface basis.

Interface card specifications

Table 2 on page 8 lists the specifications of the Lucent Metro EON system interface cards qualified by IBM @server zSeries.

Lucent Metro EON system interface cards do not support auto-negotiation of link speeds. For zSeries FICON and FCP client links, the desired link speed must be preconfigured in the WDM OTR client interface at both ends of the link.

Refer to the WDM vendor documentation for client port attenuation, distance, and link budget specifications.

Table 2 Qualified client interface card details

Card Type / Protocol	Fiber Type
LSBB, 8 Mbps to 750 Mbps, 1310 nm	
ESCON/SBCON MM	MM 50 / 62.5 um
ELSBB, 8 Mbps to 1.062 Gbps, 1310 nm	
ESCON/SBCON MM	MM 50 / 62.5 um
ETR MM	MM 50 / 62.5 um
ISC / ISC-2 / ISC-3 Compatibility Mode	SM 9 um
UBB, 8 Mbps to 2.5 Gbps, 1310 nm	
ESCON/SBCON MM	MM 50 / 62.5 um
CLO MM	MM 50 / 62.5 um
FICON 1.062 Gbps SM	SM 9 um
Fibre Channel 1.062 Gbps SM	SM 9 um
ISC-3 Peer 1 Gbps ^a SM	SM 9 um
ISC-3 Peer Mode 2 Gbps SM	SM 9 um
Gigabit Ethernet SM	SM 9 um
SM = single-mode fiber (9/125 micron) MM = multimode fiber (either 50/125 or 62.5/125 micron) MCP = Mode Conditioning Patch Cable LX = long wavelength, 1310 nm SX= short wavelength, 850 nm	

a. To support ISC-3 Peer at 1 Gbps (RPQ 8P2197), configure the WDM interface protocol as: Fibre Channel FC100 (1 Gbps).

References

For more information about zSeries connectivity, see:

<http://www.ibm.com/servers/eserver/zseries/connectivity/>

For more information about GDPS solutions, see:

- ▶ Parallel Sysplex home page:

<http://www.ibm.com/servers/eserver/zseries/psf/>

- ▶ GDPS White paper:

<http://www.ibm.com/servers/eserver/zseries/library/whitepapers/>

For more information about the IBM TotalStorage Proven program, see:

<http://www.storage.ibm.com/proven/index.html>

For more information about the Lucent Technologies Metro Enhanced Optical Network (EON)[™] System, see:

<http://www.lucent.com>

For information about other zSeries qualified WDM vendor products, see the following Redpapers and Web sites:

- ▶ *IBM 2029 Fiber Saver DWDM, SG24-5608* (withdrawn from marketing)

<http://www.redbooks.ibm.com/abstracts/sg245608.html?Open>
<http://www.ibm.com>

- ▶ *zSeries qualified WDM vendor: Adva Optical Networking, REDP-3903*

<http://www.redbooks.ibm.com/abstracts/redp3903.html?Open>
<http://www.advaoptical.com>

- ▶ *zSeries qualified WDM vendor: Cisco Systems, REDP-3905*

<http://www.redbooks.ibm.com/abstracts/redp3905.html?Open>
<http://www.cisco.com>

- ▶ *zSeries qualified WDM vendor: Nortel Networks, REDP-3904*

<http://www.redbooks.ibm.com/abstracts/redp3904.html?Open>
<http://www.nortelnetworks.com>

IBM GDPS qualification letters are available on ResourceLink:

<https://www.ibm.com/servers/resourceLink/lib03020.nsf/pages/zseriesQualifiedExtendersAndWdmProductsForGdpsSolutions?OpenDocument&pathID=>

The team that wrote this Redpaper

This Redpaper was produced by a team of specialists working at the International Technical Support Organization, Poughkeepsie Center.

Bill White is a Project Leader and Senior Networking Specialist at the International Technical Support Organization, Poughkeepsie Center.

Rav Meerwald is a Networking Systems Services Representative with IBM Australia. He has 16 years of field experience with IBM, including Networking and WDM products. He provides design, implementation, and technical assistance for WDM solutions.

Jeff Nesbitt is a zSeries Systems Services Representative with IBM Australia. He has 21 years of field experience in zSeries and S/390 servers and networking products. He also provides professional services and training to zSeries clients.

Thanks to the following people for their contributions to this project:

Connie Beuselinck
zSeries Product Planning, IBM Poughkeepsie

Casimer DeCusatis
zSeries Development Lab Connectivity Solutions, IBM Poughkeepsie

Noshir Dhondy
zSeries Product Planning, IBM Poughkeepsie

Sam Mercier
Vendor Solutions Connectivity Lab, IBM Poughkeepsie

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

This document created or updated on April 4, 2005.




Send us your comments in one of the following ways:

- ▶ Use the online **Contact us** review redbook form found at:
ibm.com/redbooks
- ▶ Send your comments in an email to:
redbook@us.ibm.com
- ▶ Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYJ Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

@server®	GDPS®	Sysplex Timer®
Enterprise Storage Server®	IBM TotalStorage Proven™	TotalStorage Proven™
ESCON®	IBM®	TotalStorage®
FICON®	Parallel Sysplex®	z/OS®
Geographically Dispersed Parallel	Redbooks (logo)  ™	zSeries®
Sysplex™	S/390®	

The following terms are trademarks of other companies:

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.